

Optimal Driving Under Traffic Signal Uncertainty[★]

Mallory E. Gaspard, Alexander Vladimirovsky^{*}

^{*} Cornell University, Ithaca, NY 14853 USA
(e-mail: { meg375, vladimirsky } @cornell.edu).

Abstract: We study driver's optimal trajectory planning under uncertainty in the duration of a traffic light's green phase. We interpret this as an optimal control problem with an objective of minimizing the expected cost based on the fuel use, discomfort from rapid velocity changes, and time to destination. Treating this in the framework of dynamic programming, we show that the probability distribution on green phase durations gives rise to a sequence of Hamilton-Jacobi-Bellman PDEs, which are then solved numerically to obtain optimal acceleration/braking policy in feedback form. Our numerical examples illustrate the approach and highlight the role of conflicting goals and uncertainty in shaping drivers' behavior.

Copyright © 2022 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Optimal control under uncertainty; dynamic programming; traffic signal advisory.

1. INTRODUCTION

Determining optimal driving behaviors under traffic signal uncertainty is of crucial importance in modern transportation applications. Recent technological advances aim to provide real-time “signal phase and timing” (SPaT) information to smartphones and on-board vehicle systems, with the goal of improving safety, fuel efficiency, and driver's comfort. While this technology will yield likely decreases in crash rates and carbon emissions (Chang et al. (2015)), the full transition to vehicles communicating directly with traffic signals will take several decades (Motavalli (2020)), and some level of uncertainty in SPaT information will remain even then due to occasional network failures and the presence of actuated or semi-actuated traffic signals (Eteifa et al. (2021)).

Modeling how drivers do (or should) utilize SPaT knowledge is a popular research area; e.g., Sun et al. (2018), Sun et al. (2020), Mahler and Vahidi (2014), Mahler and Vahidi (2012), Zhao and Zhang (2021). But typical prior approaches are based on either full SPaT information or ad hoc models of SPaT-robustness, with the emphasis placed on detailed vehicle dynamics, road models with multiple signalized intersections, traffic effects, and route-planning choices to come up with trip plans that optimize the fuel efficiency. In contrast, we consider the simplest scenario (one car with simplified dynamics on an empty road with a single traffic signal – as described in Section 2), focusing instead on a detailed model of SPaT-uncertainty due to the probability distribution over possible durations of the green light (Section 3). The driver knows the yellow light and red light durations, so they can deterministically plan the rest of their strategy once the initial green light turns yellow. They must also obey the traffic laws (red light and speed limit) and balance expected costs related to fuel

usage, aggressive acceleration / deceleration behavior, and total time to target.

We address this optimal control problem in the dynamic programming (DP) framework, as was recently done for a related *red phase duration uncertainty* problem in (Typaldos et al. (2020)). However, the dilemma in the green phase (i.e., should we try to beat the red light?) is far more affected by the available SPaT information and the duration of the transitional yellow phase. Unlike in (Typaldos et al. (2020)), where the discrete nature of DP was essential and the timestep could not be refined, our continuous DP yields a sequence of Hamilton-Jacobi-Bellman (HJB) equations, which are then solved numerically (Section 4) to determine the optimal acceleration and braking policy in feedback form. Our numerical experiments in Section 5 highlight the role of conflicting optimization goals and SPaT-uncertainty patterns in shaping drivers' behavior.

2. PROBLEM FORMULATION

Throughout this paper, we focus on the problem of a single driver aiming to optimize their route to a target through an intersection with one traffic light. We assume the vehicle is just a point mass traveling right-to-left on a flat single-lane road (an interval $[d^*, \bar{d}]$) with no other traffic: from its starting position at $d \leq \bar{d}$ with initial speed v , through the intersection at $d_\ell < d$ and towards its target at $d^* < d_\ell$. The driver varies the car's signed acceleration and must stay under the speed limit $\bar{v}$ on the way to d^* . The car may not enter the intersection (crossing d_ℓ) while the light is red, but may do so while the light is yellow¹. We further assume that the driver knows the yellow and red signal durations (D_Y, D_R).

¹ This corresponds to *permissive* yellow light laws (Beeber (2020)), where the car does not have to clear the intersection by the time the light turns red. This makes the actual intersection width and the physical car length irrelevant.

[★] Supported by the NSF DMS (awards 1645643, 1738010, 2111522) and the first author's NDSEG Fellowship.

Time Horizon. The entire planning process is split into (1) an initial green light phase I starting at $t = 0$ followed by (2) a yellow phase Y starting at $t = T_Y$, (3) a red phase R starting at $T_R = T_Y + D_Y$, and (4) an indefinite green phase G starting at $T_G = T_R + D_R$ and terminating when we reach the target at d^* . In many realistic situations, the traffic signal uncertainty will primarily rest in phase I . The limited knowledge scenario discussed in Section 3 will define a probability distribution on the *random turning yellow time* $\hat{T}_Y$, thereby leading to uncertainty in the finite time horizon of phase I .

Vehicle Dynamics. Let $\Omega = \{(d, v, r) : d \in [d^*, \bar{d}], v \in [0, \bar{v}], r \geq 0\}$. We will use $\mathbf{x} = [d \ v]^T$ to refer to the vehicle's initial configuration or a generic point in (d, v) space, switching to $\mathbf{y}(r) = [d(r) \ v(r)]^T$ to encode how position and velocity change with time. Starting from an initial configuration $\mathbf{y}(t) = \mathbf{x}$ and assuming that external forces acting on the car are negligible, the vehicle's dynamics are

$$\dot{\mathbf{y}}(r) = \mathbf{f}(\mathbf{y}(r), a(r)) = \begin{bmatrix} -v(r) \\ a(r) \end{bmatrix}, \quad \text{for } r \geq t \quad (1)$$

where the vehicle's acceleration $a(\cdot)$ is a measurable control function, $a : \mathbb{R} \rightarrow A$, and the set of available (signed) acceleration values is $A = [-\alpha, \beta]$.

Cost Function. The driver's goal on the trip from d to d^* is to select $a(\cdot)$ to address several optimization objectives: fuel consumption², discomfort from rapid acceleration/deceleration, and time to target. We aim to minimize an integral of the *running cost*

$$K(\mathbf{y}(r), a) = c_1[a]_+ + c_2 a^2 + c_3 \quad (2)$$

where $[a]_+ = \max\{a, 0\}$. The non-negative objective weights c_i 's reflect individual preferences. In principle, these can also be learned from trends in driver behavior data, as in (Butakov and Ioannou (2016)).

Deterministic Control Problem. We begin by discussing the fully deterministic problem that starts after the light turns yellow (i.e., the optimal driving for $t \geq T_Y$). In the usual style of dynamic programming, we consider the remaining phases of the traffic light backwards in time.

We first consider the last green phase of unlimited duration, G , which is best described as an *exit-time problem*, terminating when the vehicle reaches its target at the time $T^* = \min\{r \geq T_G \mid \mathbf{y}(r) = d^*\}$. The cost-to-go depends on our initial configuration and the chosen control $\mathbf{a}(\cdot)$, which together determine the remaining time to target $(T^* - t)$. We thus define the control-specific cost

$$\mathcal{J}_G(\mathbf{x}, a(\cdot)) = \int_t^{T^*} K(\mathbf{y}(r), a(r)) \, dr, \quad (3)$$

starting from $\mathbf{x} = \mathbf{y}(t)$ with $t \geq T_G$. By a standard argument (Bardi and Capuzzo-Dolcetta (1997)), the value function defined as

$$q(\mathbf{x}) = \inf_{a(\cdot)} \{\mathcal{J}_G(\mathbf{x}, a(\cdot))\} \quad (4)$$

can be recovered as a viscosity solution of a stationary Hamilton-Jacobi-Bellman (HJB) equation

$$0 = \min_{a \in A} \{K(\mathbf{x}, a) + \nabla q \cdot \mathbf{f}(\mathbf{x}, a)\} \quad (5)$$

with boundary conditions $q(d^*, v) = 0, \forall v \in [0, \bar{v}]$.

During the other two deterministic stages (Y and R), we define the control-specific cost as

$$\mathcal{J}(\mathbf{x}, t, a(\cdot)) = \int_t^{T_G} K(\mathbf{y}(r), a(r)) \, dr + q(\mathbf{y}(T_G)), \quad (6)$$

starting from $\mathbf{x}$ at the time $t \in [T_Y, T_G]$. The value function

$$u(\mathbf{x}, t) = \inf_{a(\cdot)} \{\mathcal{J}(\mathbf{x}, t, a(\cdot))\} \quad (7)$$

is now time-dependent to account for the fixed time-horizon. The state constraints detailed below can often make this value function discontinuous during the Y phase. Nevertheless, u can still be recovered as a discontinuous viscosity solution (Bardi and Capuzzo-Dolcetta, 1997, Chapter 5) of the time-dependent HJB equation

$$0 = u_t + \min_{a \in A} \{K(\mathbf{x}, a) + \nabla u \cdot \mathbf{f}(\mathbf{x}, a)\} \quad (8)$$

subject to the terminal condition $u(\mathbf{x}, T_G) = q(\mathbf{x})$.

State Constraints. The car should never enter the intersection during the R phase³. This makes it natural to consider a space-time obstacle $\mathcal{I} = \{(d, v, t) \mid d = d_\ell, 0 \leq v \leq \bar{v}, T_R \leq t < T_G\}$, but we go further and forbid the car from taking any (d, v) from which it is impossible to come to a complete stop before reaching d_ℓ . (This guarantees that the car will not violate the rules even if the signal fails to switch to green light at the usual time T_G .) The minimal stopping distance depends on the current velocity v and the maximum deceleration rate α , yielding the curve of critical positions (the “parabola of last resort”) $d_\alpha(v) = d_\ell + \frac{v^2}{2\alpha}$ and the disallowed region

$$\mathcal{I}_R(t) = \{(d, v) \mid d_\ell \leq d < d_\alpha(v), 0 \leq v \leq \bar{v}\}$$

enforced for $t \in [T_R, T_G]$. Since we are forced to use the maximum deceleration $a = -\alpha$ on the parabola of last resort, this makes it easy to compute the value function on it. Let $s_G = T_G - t$. It would take the vehicle $s_\alpha = v/\alpha$ to come to a full stop, but the actual braking time might be shorter, $\hat{s}_\alpha = \min\{s_\alpha, s_G\}$. The value function along $(d_\alpha(v), v, t)$ is

$$C_\alpha = \hat{s}_\alpha c_2 \alpha^2 + c_3 s_G + q(\mathbf{y}(T_G)). \quad (9)$$

During the Y phase, the car is allowed to have $d \leq d_\alpha(v)$ as long as it has enough time to enter the intersection by the time T_R . The *largest* d starting from which this is possible can be found by assuming that we use the maximum allowable acceleration (but without violating the speed limit) until reaching d_ℓ . Denote the remaining yellow time as $s_R = T_R - t$, the time that it would take to reach the speed limit as $s_\beta = \frac{\bar{v}-v}{\beta}$, and the lowest starting speed from which the speed limit would be reached before T_R as $v_c = \bar{v} - \beta s_R$. For fixed v and $t \in [T_Y, T_R]$, the largest allowed d is then

² We approximate fuel consumption as the amount of fuel let into the engine. We assume that fuel enters the engine at a rate proportional to the car's positive acceleration (i.e., how far down the driver is pressing the gas pedal) (Bennett (2016)).

³ For the sake of brevity, we omit detailed algebraic derivations in favor of geometric discussion emphasizing the physical intuition and main results. Further details of the numerical implementation are discussed in Section 4.

$$d_\beta(v, t) = \begin{cases} d_\ell + v s_R + \frac{\beta}{2} s_R^2 & v \leq v_c; \\ d_\ell + v s_\beta + \frac{\beta}{2} s_\beta^2 + \bar{v}(s_R - s_\beta), & v \geq v_c. \end{cases} \quad (10)$$

The new disallowed region

$$\mathcal{I}_Y(t) = \{(d, v) \mid d_\beta(v, t) \leq d \leq d_\alpha(v), v \in [0, \bar{v}]\}$$

is enforced for $t \in [T_Y, T_R]$ and illustrated in Fig. 1. When $d_\beta(v, t) < d_\alpha(v)$, the value function along $(d_\beta(v, t), v, t)$ is

$$C_\beta = (c_1\beta + c_2\beta^2)\hat{s}_\beta + c_3s_R + u(\mathbf{y}(T_R), T_R) \quad (11)$$

where $\hat{s}_\beta = \min\{s_\beta, s_R\}$. We assume that the yellow phase duration D_Y is long enough so that, starting from every (d, v) , it is possible to either enter the intersection before T_R or come to a full stop before the intersection; i.e., $\mathcal{I}_Y(T_Y) = \emptyset$. Whenever $d_\alpha(v) \leq d_\beta(v, t)$, the latter starting position allows either option, and if speeding toward the intersection yields a lower cost, then the value function will be discontinuous at $(d_\beta(v, t), v, t)$.




Fig. 1. Two snapshots of the time-dependent piecewise boundary during phase Y . Allowed (d, v) configurations are shown in green, and $\mathcal{I}_Y(t)$ is shown in white. For the starting configurations $(d_\beta(v, t), v)$, if speeding through the light before $t = T_R$ yields a lower cost than waiting for $T = T_G$, then the value function u is discontinuous along the orange-dashed lines. As $t \rightarrow T_R$, $d_\beta \rightarrow d_\ell$, and $\mathcal{I}_Y(t) \rightarrow \mathcal{I}_R$.

3. OPTIMAL DRIVING UNDER T_Y UNCERTAINTY

Optimal planning in the initial green phase is quite similar if the duration ($D > 0$) and the starting time of that phase are fully known. For a car that has been driving toward the intersection ever since this light turned green ξ seconds before the current time $t = 0$, the natural planning horizon is $T = D - \xi$. The light will then turn yellow, and the remaining cost-to-go will be $\delta(\mathbf{x}) = u(\mathbf{x}, T_Y)$ already computed in Section 2. The new value function $w(\mathbf{x}, t)$ can be obtained by solving (8) with $t \in [0, T]$ and the terminal condition $w(\mathbf{x}, T) = \delta(\mathbf{x})$. But in this section we focus on scenarios where D is specified only probabilistically while the durations of the subsequent phases (D_Y and D_R) are known and fixed. For a simple example, consider a city where all traffic light signals have green phase durations of either D_1 or $D_2 > D_1$ seconds, and the fractions of signals in these respective categories are p_1 and $p_2 = 1 - p_1$. Even though the driver is aware of the city-wide statistics, they may not know the category of this particular signal. So, the remaining green time is a random variable $\tilde{T}_Y$ with possible values $T_i = D_i - \xi$ and the corresponding probabilities p_i , ($i = 1, 2$). We are thus selecting a control

$\mathbf{a}(\cdot)$ to be used until the light turns yellow with the goal of minimizing

$$\mathcal{C}(\mathbf{x}, t, \mathbf{a}(\cdot)) = \mathbb{E} \left\{ \int_t^{\tilde{T}_Y} K(\mathbf{y}(r), \mathbf{a}(r)) dr + \delta(\mathbf{y}(\tilde{T}_Y)) \right\}. \quad (12)$$

This is an example of “initial uncertainty” (Qi et al. (2021)), with the problem becoming fully deterministic at the time T_1 : the light will either turn yellow then or we will immediately know that $\tilde{T}_Y = T_2$. Thus, starting from any $t < T_1$,

$$\begin{aligned} \mathcal{C}(\mathbf{x}, t, \mathbf{a}(\cdot)) &= \sum_{i=1}^2 p_i \left(\int_t^{T_i} K(\mathbf{y}(r), \mathbf{a}(r)) dr + \delta(\mathbf{y}(T_i)) \right) \\ &= \int_t^{T_1} K(\mathbf{y}(r), \mathbf{a}(r)) dr + p_1 \delta(\mathbf{y}(T_1)) + \\ &\quad (1 - p_1) \left(\int_{T_1}^{T_2} K(\mathbf{y}(r), \mathbf{a}(r)) dr + \delta(\mathbf{y}(T_2)) \right). \end{aligned}$$

This makes it natural to treat the planning on $[0, T_1]$ and $[T_1, T_2]$ as separate optimization problems, defining a pair of value functions $w^1(\mathbf{x}, t)$ and $w^2(\mathbf{x}, t)$ for the respective time intervals. Each of these will solve the same PDE (8), but with different terminal conditions

$$w^1(\mathbf{x}, T_1) = p_1 \delta(\mathbf{x}) + (1 - p_1) w^2(\mathbf{x}, T_1); \quad w^2(\mathbf{x}, T_2) = \delta(\mathbf{x}).$$

Generalizing this to n possible green light durations $D_1 < \dots < D_n$ with possible remaining green time intervals $T_i = D_i - \xi$ and $\mathbb{P}(\tilde{T}_Y = T_i) = p_i$ for $i = 1, \dots, n$, we introduce n different value functions $w^i(\mathbf{x}, t)$, each defined on its own time interval $[T_{i-1}, T_i]$, where for notational convenience we take $T_0 = 0$. All of these satisfy (8) with the terminal conditions

$$w^i(\mathbf{x}, T_i) = \hat{p}_i \delta(\mathbf{x}) + (1 - \hat{p}_i) w^{i+1}(\mathbf{x}, T_i), \quad (13)$$

where $\hat{p}_i$ is the conditional probability

$$\hat{p}_i = \mathbb{P}(\tilde{T}_Y = T_i \mid \tilde{T}_Y > T_{i-1}) = p_i / \left(\sum_{j=i}^n p_j \right).$$

Since $\hat{p}_n = 1$, the problem is deterministic on $[T_{n-1}, T_n]$. For $i < n$, each w^i depends only on the next w^{i+1} , and all value functions can be found in a single sweep backward in time from $t = T_n$ to $t = T_0$.

4. NUMERICAL IMPLEMENTATION

4.1 Numerical Methods for HJB Equation

Overall, our numerical approach to the “uncertain green phase duration” problem consists of the following three stages:

- (1) Solve the stationary HJB (5) for the last phase G .
- (2) Solve HJB (8) for the deterministic phases R and Y .
- (3) Solve a sequence of HJBs (8) for the uncertain phase I .

The numerical results from each stage are used in determining the terminal conditions for the following stage. Our algorithms are based on a semi-Lagrangian (SL) discretization (Falcone and Ferretti (2014)) on a Cartesian grid over a (d, v, t) domain with $(N_d + 1), (N_v + 1)$, and

$(N_t + 1)$ gridpoints along the respective dimensions. For a pre-specified N_v , we select the discretization parameters

$$\Delta v = \frac{\bar{v}}{N_v}, \quad \Delta t = \frac{\Delta v}{\max(\alpha, \beta)}, \quad \Delta d = \bar{v} \Delta t. \quad (14)$$

The physical position, velocity, and time at a node (i, j, k) are then $(\mathbf{x}_{ij}, t_k) = (d_i, v_j, t_k) = (d^* + i\Delta d, j\Delta v, k\Delta t)$ for $i = 0, \dots, N_d$, $j = 0, \dots, N_v$, and $k = 0, \dots, N_t$. The usual first-order SL discretization is based on assuming that a fixed control a is used starting from $\mathbf{x}_{ij}$ for a small time τ . The resulting new state $\tilde{\mathbf{x}}_{ij}^a = (\tilde{d}_i^a, \tilde{v}_j^a)$ is usually approximated numerically (e.g., as $\tilde{\mathbf{x}}_{ij}^a \approx \mathbf{x}_{ij} + \tau \mathbf{f}(\mathbf{x}_{ij}, a)$), but our simplified dynamics (1) allows computing $\tilde{\mathbf{x}}_{ij}^a$ analytically. Moreover, (14) guarantees that $\tilde{d}_i^a \in [d_{i-1}, d_i]$ and $|\tilde{v}_j^a - v_j| \leq \Delta v$ for any $\tau \leq \Delta t$, which makes it easier to interpolate the value function at $\tilde{\mathbf{x}}_{ij}^a$. We also use the Golden Section Search (GSS) algorithm whenever we need to find the optimal a numerically.

Stage 1: Stationary HJB Solve In phase G , there is never any incentive to decelerate and we only consider control values $a \in [0, \beta]$. The PDE (5) is stationary, and the solution $q(\mathbf{x}_{ij})$ is approximated by a grid function Q_{ij} satisfying

$$\begin{aligned} Q_{ij} &= \min_{a \in [0, \beta]} \{ \tau K(\mathbf{x}_{ij}, a) + Q(\tilde{\mathbf{x}}_{ij}^a) \}, \quad i > 0, j < N_v; \\ Q_{iN_v} &= \tau K(\mathbf{x}_{iN_v}, 0) + Q(\tilde{\mathbf{x}}_{iN_v}^0), \quad \forall i; \\ Q_{0j} &= 0, \quad \forall j; \end{aligned} \quad (15)$$

where the last two lines encode the need to stay under the speed limit and the boundary conditions respectively. Since $\tilde{\mathbf{x}}_{ij}^a$ is usually not a gridpoint, its value Q is recovered via bilinear interpolation. I.e.,

$$Q(\tilde{\mathbf{x}}_{ij}^a) = \gamma_1(a)Q_{i,j} + \gamma_2(a)Q_{i,j+1} + \gamma_3(a)Q_{i-1,j+1} + \gamma_4(a)Q_{i-1,j},$$

with the nonnegative bilinear coefficients adding up to 1 and $\gamma_1(a) < 1$ for all a values. The coupled system (15) can be solved through *value iterations*, but the dependence of the right hand side on Q_{ij} slows down the convergence considerably even when using a Gauss-Seidel relaxation. Instead, we replace the first equation in (15) with an equivalent

$$Q_{ij} = \min_{a \in [0, \beta]} \left\{ \frac{1}{1 - \gamma_1(a)} \left(\tau K(\mathbf{x}_{ij}, a) + \gamma_2(a)Q_{i,j+1} + \gamma_3(a)Q_{i-1,j+1} + \gamma_4(a)Q_{i-1,j} \right) \right\}. \quad (16)$$

This effectively decouples the system, and we can now solve it in one sweep, looping through $i = 1, \dots, N_d$; $j = N_v, \dots, 0$.

Stage 2: Deterministic Y-R Phase HJB Solve The SL discretization for the time-dependent PDE (8) is similar, but with several subtleties, which we describe below only briefly due to the space constraints. The grid-approximation of the value function must satisfy

$$U_{ij}^k = \min_{a \in A_{ij}} \{ K(\mathbf{x}_{ij}, a) \tau + U(\tilde{\mathbf{x}}_{ij}^a, t + \tau) \} \quad (17)$$

for all i, j and all $k < N_t$, which can be solved backwards in time from the terminal conditions $U_{ij}^{N_t} = Q_{ij}$. The control set is based on the speed constraints: $A_{i0} = [0, \beta]$, $A_{iN_v} = [-\alpha, 0]$, and $A_{ij} = [-\alpha, \beta]$ for all other j values. On most

of the domain we use $\tau = \Delta t$, and the value at $(\tilde{\mathbf{x}}_{ij}^a, t + \tau)$ is obtained via bilinear interpolation in (d, v) . But we reduce τ wherever it is needed to respect additional state constraints due to the traffic light, in which case we employ *cut-cells* in (d, v) and additional linear interpolation in t .

In the red phase (with $t \in [T_R, T_G]$), the car has a disallowed region $\mathcal{I}_R$. For $\mathbf{x}_{ij}$ just right of the line $d = d_\alpha(v)$, our τ is decreased adaptively, to ensure that $\tilde{d}_{ij}^a \geq d_\alpha(\tilde{v}_{ij}^a)$. All grid cells intersected by this line are treated as cut-cells, with values of U on the parabolic boundary computed by formula (9).

Extending the solution into the yellow phase (for $t \in [T_Y, T_R]$) presents two additional complications: the time-dependent disallowed region $\mathcal{I}_Y(t)$ (see Fig. 1) and a possible discontinuity of $u(\mathbf{x}, t)$ when $d = d_\beta(v, t)$. We adopt the following two-pass procedure to handle these challenges separately. First, we solve (17) for all $d \geq d_\alpha(v)$ exactly as we did in phase R ; i.e., treating $d = d_\alpha(v)$ as a boundary with the boundary conditions specified by (9). This computes the best cost attainable by waiting out the red light. But for some starting configurations, it is also possible to accelerate enough and cross the intersection before the light turns red. So, we then re-solve (17) for all $d \leq d_\beta(v, t)$ without treating $d = d_\alpha(v)$ as a boundary, with grid values updated only if they are smaller than those obtained in the first pass. To avoid interpolating across the discontinuity, in this second pass we use cut-cells just left of $d = d_\beta(v, t)$ with the boundary conditions specified by (11).

Stage 3: Uncertain Initial Phase I Given a set of possible times when the light could turn yellow $T_1, \dots, T_n$, each conditional value function $w^i(\mathbf{x}, t)$ satisfies PDE (8) on $t \in [T_{i-1}, T_i]$. Thus, we solve them sequentially (from w^n to w^1) using the same SL-discretization (17) but with the terminal conditions (13). We note that the latter uses $\delta(\mathbf{x}) = u(\mathbf{x}, T_Y)$ obtained at the end of Stage 2. Once this function is computed, Stage 3 can be repeated as needed (for different probability distributions of $\tilde{T}_Y$) without repeating the computations in Stages 1 and 2. We also note that the timestep in stage I is always $\tau = \Delta t$ and cut-cells are not needed since $\mathcal{I}_Y(T_Y) = \emptyset$.

4.2 Optimal Trajectory Tracing

Once the value function is computed, it is easy to find an approximately optimal trajectory for every starting configuration (d, v, t) . For $t \in [T_Y, T_G]$, the optimal control value a_* is stored at every gridpoint while solving (17); at all other points, we approximate $a_*(\mathbf{x}, t)$ via tri-linear interpolation. When tracing a trajectory, we use this approximately optimal feedback control to solve (1) until the time T_G . At that point, we switch to the approximately optimal feedback control obtained from (15) and continue until reaching the target d^* .

If starting in the initial green phase of uncertain duration from some $t \in [T_{i-1}, T_i]$, we again follow the feedback control obtained when solving for w^i . If the light turns yellow at the time T_i , we switch to a deterministic setting with the feedback control based on u . Otherwise, (i.e., if it stays green at T_i), we continue with w^{i+1} .

4.3 Approximating the $\mathcal{J}_3$ -Constrained Pareto Front

Our running cost K defined in (2) can be viewed as a linear combination of constituent running costs describing separately the fuel usage, discomfort from rapid speed changes, and time to target; i.e., $K_1 = \gamma[a]_+$, $K_2 = a^2$, and $K_3 = 1$, with the coefficients (c_1, c_2, c_3) reflecting the relative importance of these to the driver. Cumulative constituent costs $(\mathcal{J}_1, \mathcal{J}_2, \mathcal{J}_3)$ can be computed by integrating the corresponding K_i 's along the trajectory. Rational tradeoffs among these objectives can be explored by varying ratios of c_i 's. To evaluate the fuel-discomfort tradeoffs among trajectories reaching the target in at most $T^\dagger$ seconds, we fix a ratio $\frac{c_1}{c_2}$ and find c_3 such that $\mathcal{J}_3 = T^\dagger$ for the K -optimal trajectory recovered from the HJB PDE. Repeating this process for a range of $\frac{c_1}{c_2}$ values, we obtain a $\mathcal{J}_3$ -constrained Pareto Front for $(\mathcal{J}_1, \mathcal{J}_2)$.

5. NUMERICAL EXPERIMENTS

For all experiments⁴, we use realistic parameter values $\bar{v} = 45 \text{ mph} \approx 20.12 \text{ m/s}$, $-\alpha = -3.8 \text{ m/s}^2$, $\beta = 3.8 \text{ m/s}^2$, $D_Y = 3 \text{ s}$, $D_R = 60 \text{ s}$. We conduct all planning on a 200-meter-long road segment with the traffic light in the middle; i.e., $\bar{d} = 100 \text{ m}$, $d_\ell = 0$, and $d^* = -100 \text{ m}$. We use a 345×181 grid in (d, v) space, with Δt specified by (14).

Mirroring the previous sections, we start with the planning in the later (deterministic) phases before considering the effects of uncertainty in the initial phase I .

Example 1: Green phase G only. The planning in this final green phase is particularly easy to interpret since there are no constraints other than the speed limit. With our simplified dynamics, the method of characteristics used on PDE (5) shows that any optimal trajectory will have at most three stages: (1) using the maximum acceleration $a = \beta$, (2) linearly decreasing acceleration until $a = 0$, and (3) coasting until we reach the target. We illustrate this for a specific starting $(d, v) = (80, 0)$ in Fig. 2A. Starting from the same (d, v) , Fig. 2B shows the rational (fuel usage / acceleration discomfort) tradeoffs for 3 different constraint levels on the time-to-target.




Fig. 2. Example 1 optimal trajectory and fuel-discomfort tradeoffs starting from $(d, v) = (80, 0)$. (A): Optimal $d(t)$, $v(t)$, and $a(t)$ when $(c_1, c_2, c_3) = (0.025, 0.025, 0.95)$. The individual costs along the trajectory are $\mathcal{J}_1 = 20.11$, $\mathcal{J}_2 = 66.82$ and $\mathcal{J}_3 = 11.73\text{s}$. (B): $\mathcal{J}_3$ -constrained Pareto fronts for $(\mathcal{J}_1, \mathcal{J}_2)$ with $\mathcal{J}_3 = 25\text{s}$, 35s , and 45s .

⁴ In the interest of computational reproducibility, we provide the full source code, additional figures, and movies for all examples at https://eikonal-equation.github.io/Traffic_Light_Uncertainty/

Example 2: Phases R and G . Figure 3 shows what happens when the driver starts planning as the light turns red at $t = T_R$. Beyond the traffic light, for $d < d_\ell$, the optimal policy is exactly the same as in phase G . But on the other side, the car cannot cross the parabolic boundary $d_\alpha(v) = d$ until $T_G = T_R + 60\text{s}$. This requires aggressive braking to the right of this parabola, though the need for harsh braking diminishes as $t \rightarrow T_G$ and $d \rightarrow \bar{d}$ since it is optimal to reach $d_\alpha(v) = d$ with non-zero speed just as the light turns green.




Fig. 3. Example 2 feedback controls and optimal trajectory when $(c_1, c_2, c_3) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. (A): Feedback controls at $t = T_R$ with $\mathcal{I}_R$ shown in white. (B): Optimal $d(t)$, $v(t)$ and $a(t)$ for starting $(d, v, t) = (80, 15, T_R)$. It is optimal for the driver to brake, stop, and wait until they can accelerate to arrive at $d_\alpha(v)$ at T_G with $v > 0$.

Example 3: Deterministic phases Y, R , and G . In the yellow phase, some of the starting configurations (with $d \leq d_\beta(v, t)$) allow beating the red light if the driver is ready to accelerate aggressively. This leads to a time-dependent disallowed set $\mathcal{I}_Y(t)$ (see Fig. 4) and a possible discontinuity in the value function when $d = d_\beta(v, t)$. Fig. 5 shows that even a slight shift in a starting position over this discontinuity will result in a drastically different optimal trajectory and cumulative cost.




Fig. 4. Example 3 feedback controls for $(c_1, c_2, c_3) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ at (A): $t = T_Y + 1\text{s}$ and (B): $t = T_Y + 2\text{s}$. Both plots are the control-heatmap equivalents to the piecewise boundary schematics shown in Fig. 1.

Example 4: Uncertain yellow change, $\tilde{T}_Y \in \{T_1, T_2\}$. We now start planning at the time $t = 0$ in the initial green phase I , with the remaining time until yellow $\tilde{T}_Y$ taking values $T_1 = 2\text{s}$ or $T_2 = 6\text{s}$ with equal probability ($p_1 = p_2 = 1/2$). We first assume that the driver values all three objectives equally ($c_1 = c_2 = c_3 = \frac{1}{3}$); the resulting optimal feedback control for $t = 0$ is shown in Fig. 6A. We focus on the starting configuration $(d, v) = (94, 0.85)$, from which it is possible to beat the traffic light if $\tilde{T}_Y = T_2$ but not if $\tilde{T}_Y = T_1$. The optimal trajectory (shown in Fig. 6B) branches at the time $t = T_1$, when we find out the true value of $\tilde{T}_Y$. We note that the “optimal under uncertainty”




Fig. 5. Example 3 optimal trajectories starting from (A): $(d, v, t) = (43, 10, T_Y)$ and (B): $(d, v, t) = (48, 10, T_Y)$. The former yields $(\mathcal{J}_1, \mathcal{J}_2, \mathcal{J}_3) = (6.82, 21.32, 8.9)$ and the overall cost $\mathcal{J} = 12.35$. The latter results in $(\mathcal{J}_1, \mathcal{J}_2, \mathcal{J}_3) = (6.65, 14.16, 80.95)$ and a much larger $\mathcal{J} = 33.94$.

control used for $t \in [0, T_1]$ would not be optimal for either deterministic scenario: we would accelerate far less (if at all) with $T_Y = T_1$ and far more with $T_Y = T_2$.

But just because one can beat the traffic light with $\tilde{T}_Y = T_2$, it does not mean that it is always optimal to do so. Indeed, if this light duration is sufficiently unlikely, the initial acceleration on $[0, T_1]$ to preserve both options is no longer worthwhile. We demonstrate this for $(p_1, p_2) = (0.95, 0.05)$ in Fig. 7. (Note that there is still branching at $t = T_1$ since at that point we discover the time $(\tilde{T}_Y + D_Y + D_R)$ by which we need to reach the parabolic boundary as the light turns green again.) Not surprisingly, a similar decision not to rush can also result from a difference in driver's priorities; e.g., a high enough c_2 will make rapid acceleration unattractive even if T_2 is fairly likely. In Fig. 8 we demonstrate this for $(p_1, p_2) = (0.5, 0.5)$ and $(c_1, c_2, c_3) = (0.15, 0.75, 0.1)$.




Fig. 6. Example 4 feedback controls and trajectories for starting $(d, v) = (94, 0.85)$ when $(p_1, p_2) = (0.5, 0.5)$ and $(c_1, c_2, c_3) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. (A): Feedback controls at $t = 0$. The “X” indicates the vehicle's starting point. (B): Optimal $d(t)$, $v(t)$ and $a(t)$. Vertical lines corresponding to the possible turning yellow, red, and green times are shown in their respective colors.

Example 5: Uncertain yellow change, $\tilde{T}_Y \in \{T_1, T_2, T_3\}$. In our final example, we use $(T_1, T_2, T_3) = (2s, 4s, 6s)$ with $(p_1, p_2, p_3) = (0.25, 0.25, 0.5)$ and $c_1 = c_2 = c_3 = \frac{1}{3}$ starting from $(d, v) = (68, 5)$. The optimal control (shown in Fig. 9) has three branches: if the light stays green at $t = T_1$, it becomes clear that we can beat the red light, but how much acceleration will be needed to do so will be revealed at $t = T_2$.

6. CONCLUSIONS

We provide a framework for determining optimal driving strategies in the face of initial traffic light uncertainty and




Fig. 7. Example 4 feedback controls and trajectories for starting $(d, v) = (94, 0.85)$ when $(p_1, p_2) = (0.95, 0.05)$, and $(c_1, c_2, c_3) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. (A): Feedback controls at $t = 0$. The “X” indicates the vehicle's starting point. (B): Optimal $d(t)$, $v(t)$ and $a(t)$. Vertical lines corresponding to the possible turning yellow, red, and green times are shown in their respective colors.




Fig. 8. Example 4 feedback controls and trajectories for starting $(d, v) = (94, 0.85)$ when $(p_1, p_2) = (0.5, 0.5)$, and $(c_1, c_2, c_3) = (0.15, 0.75, 0.1)$. (A): Feedback controls at $t = 0$. The “X” indicates the vehicle's starting point. (B): Optimal $d(t)$, $v(t)$ and $a(t)$. Vertical lines corresponding to the possible turning yellow, red, and green times are shown in their respective colors.




Fig. 9. Example 5 feedback controls and trajectories for starting $(d, v) = (68, 5)$ when $(p_1, p_2, p_3) = (0.25, 0.25, 0.5)$, and $(c_1, c_2, c_3) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. (A): Feedback controls at $t = 0$. The “X” indicates the vehicle's starting point. (B): Optimal $d(t)$, $v(t)$, and $a(t)$. The (T_1, T_2, T_3) and their corresponding T_R and T_G are marked by the yellow, red, and green vertical lines respectively.

competing optimization objectives. While our basic setup is intentionally simple, the proposed optimization under uncertainty approach is much broader. We hope that it will be useful in modifying prior models with detailed vehicle dynamics (Sun et al. (2020)), route selection on complex road networks with many signalized intersections (Mahler and Vahidi (2014)), and game-theoretic traffic effects due to independent decision making of multiple drivers (Huang et al. (2019)). A similar approach will also be useful in treating other uncertainty models with more general (e.g., continuous) phase-duration distributions as

well as pedestrian-actuated signal timing changes (Eteifa et al. (2021)).

Acknowledgements: The authors are grateful to A. Nellis, J. van Hook, and N. Do for their preliminary work on the Phase R problem during the summer 2018 REU Program.

REFERENCES

- Bardi, M. and Capuzzo-Dolcetta, I. (1997). *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*. Birkhauser.
- Beeber, J. (2020). An explanation of Mats Jarlstrom’s extended kinematic equation. *ITE Journal*, 90(3), 34–38.
- Bennett, S. (2016). *Medium/Heavy Duty Truck Engines, Fuel & Computerized Management Systems*. Cengage Learning.
- Butakov, V.A. and Ioannou, P. (2016). Personalized driver assistance for signalized intersections using v2i communication. *IEEE Trans. Intell. Transp. Syst.*, 17(7), 1910–1919.
- Chang, J. et al. (2015). Estimated benefits of connected vehicle applications: dynamic mobility applications, AERIS, V2I safety, and road weather management applications. Technical Report FHWA-JPO-15-255, US Dept. of Transportation.
- Eteifa, S., Rakha, H.A., and Eldardiry, H. (2021). Predicting Coordinated Actuated Traffic Signal Change Times using Long Short-Term Memory Neural Networks. *Transp. Res. Rec.*, 03611981211000748.
- Falcone, M. and Ferretti, R. (2014). *Semi-Lagrangian approximation schemes for linear and Hamilton-Jacobi equations*, volume 133. SIAM.
- Huang, K., Di, X., Du, Q., and Chen, X. (2019). A game-theoretic framework for autonomous vehicles velocity control: Bridging microscopic differential games and macroscopic mean field games. *arXiv preprint arXiv:1903.06053*.
- Mahler, G. and Vahidi, A. (2012). Reducing idling at red lights based on probabilistic prediction of traffic signal timings. In *2012 American Control Conference (ACC)*, 6557–6562. IEEE.
- Mahler, G. and Vahidi, A. (2014). An optimal velocity-planning scheme for vehicle energy efficiency through probabilistic prediction of traffic-signal timing. *IEEE Trans. Intell. Transp. Syst.*, 15(6), 2516–2523.
- Motavalli, J. (2020). It’s time for smart traffic lights. Autoweek. URL <https://www.autoweek.com/news/a34498280/its-time-for-smart-traffic-lights/>.
- Qi, D., Dhillon, A., and Vladimirsky, A. (2021). Optimality and robustness in path-planning under initial uncertainty. preprint: <https://arxiv.org/abs/2106.11405>.
- Sun, C., Guanetti, J., Borrelli, F., and Moura, S.J. (2020). Optimal eco-driving control of connected and autonomous vehicles through signalized intersections. *IEEE Internet Things J.*, 7(5), 3759–3773.
- Sun, C., Shen, X., and Moura, S. (2018). Robust optimal eco-driving control with uncertain traffic signal timing. In *2018 Annual American Control Conference (ACC)*, 5548–5553. IEEE.
- Typaldos, P., Kalogianni, I., Mountakis, K.S., Papamichail, I., and Papageorgiou, M. (2020). Vehicle trajectory specification in presence of traffic lights with known or uncertain switching times. *Trans. Res. Rec.*, 2674(8), 53–66.
- Zhao, S. and Zhang, K. (2021). Online predictive connected and automated eco-driving on signalized arterials considering traffic control devices and road geometry constraints under uncertain traffic conditions. *Transportation research part B: methodological*, 145, 80–117.